

Cloudera Management Console on on premises Diagnostics

Date published: 2023-12-16

Date modified: 2026-08-26

CLOUDERA

Legal Notice

© Cloudera Inc. 2026. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

- Managing Cloudera on premises diagnostic data..... 4**
 - Options for generating the Cloudera on premises diagnostic data..... 4
 - Cloudera Data Services on premises diagnostic data..... 6
 - Collecting Cloudera on premises diagnostics data..... 7

Managing Cloudera on premises diagnostic data

Cloudera on premises enables you to generate and download diagnostic data associated with the various services and workloads for troubleshooting purposes. You can perform these tasks from the Diagnostics tab of the Cloudera Management Console.

Cloudera on premises uses [fluentd](#) to collect logs from the Cloudera Management Console services, and the Cloudera AI, Cloudera Data Warehouse, and Cloudera Data Engineering workloads. If Ozone is configured to store logs from the workloads, the information collected by fluentd is stored on Ozone. Otherwise, the data is moved through a fluentd aggregator in the Cloudera Management Console to an OpenShift Container Platform persistent volume where all the Cloudera Management Console logs are also stored.

When an administrator selects the option to generate the diagnostic bundle from the Cloudera Management Console > Diagnostics page, Cloudera Management Console generates a bundle containing the following details after querying the relevant components and services:

- Logs from the OpenShift Container Platform persistent volume
- Cloudera AI, Cloudera Data Warehouse, and Cloudera Data Engineering logs stored on Ozone
- Kubernetes deployment information such as events, logs, details of the OpenShift Container Platform pods and so on.
- Metrics for monitoring such as usage information and so on.
- Cloudera Management Console information such as version, UUID, and so on.

Options for generating the Cloudera on premises diagnostic data

You can specify various options on the Collect and Send Diagnostic Data pop-up window to generate Cloudera on premises diagnostic data based on your specific requirements of duration, scope, and size.

Collect New Bundle

Select the duration for which you want to generate the diagnostic data.

You can select from the following options:

- No Time Limit: Select this value if you want all the diagnostic data to be generated from the time you deployed Cloudera on premises.
- Recent Time Range: Select a specified time range, in hours, from the drop-down list.

Time Range

The duration for which you want to collect the diagnostic data. You can either select from a list of durations or specify a custom date and time range.

☐ No Time Limit
 ☒ Recent Time Range
 ☐ Custom Time Range

Recent Time Range

Last 1 day

Last 1 day

Last 2 days

Last 7 days

Last 14 days

and workloads for which you want to collect the diagnostic data.

☐ Events and Control Plane
 ☐ Environment
 ☐ Custom

Cloudera will collect diagnostic data on all environments and the control plane.

- Custom Time Range: Select the start date and end date between which you want to generate the diagnostic data.

Time Range

The duration for which you want to collect the diagnostic data. You can either select from a list of durations or specify a custom date and time range.

☐ No Time Limit ☐ Recent Time Range ☒ Custom Time Range

Start date

dd / mm / yyyy

Start Time

-- : -- AM

End Time

-- : -- AM

« < Aug 2026 > »

Mo	Tu	We	Th	Fr	Sa	Su
27	28	29	30	31	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

Scope

Select the services for which you want to collect the diagnostic data.

You can select environment- and workload-related options from the following:

- All Environments and Control Plane: Select this option to generate the diagnostic data for all the environments and workloads deployed on Cloudera on premises.

Scope

The components and workloads for which you want to collect the diagnostic data.

☒ All Environments and Control Plane ☐ Environment ☐ Custom

Cloudera will collect diagnostic data on all environments and the control plane.

- Environment: Select one or more environments from the drop-down list to generate the diagnostic data. You can also use regular expressions to filter the name of environments for which you want to generate the diagnostic data.

Storage

You can specify the size of the diagnostic data to collect. The default size of the data after compression is 500 MB. If you want to change the default size, set the BUNDLE_SIZE_LIMIT_MB environment variable on the OpenShift Container Platform pods on your Cloudera on premises deployment.

Note: The actual size of the generated diagnostic bundle might vary from the size that you specify.

While Cloudera can exclude namespace pod logs, control plane archive logs, and other archive logs from the diagnostic bundle to be consistent with the specified size, Cloudera must always include certain type of diagnostic data in the bundle. Because the size of the diagnostic data that is always included in a bundle is not constant, the overall size of the generated diagnostic bundle might vary.

Cloudera Data Services on premises diagnostic data

You can collect and download Cloudera Data Services on premises diagnostic data using the Bundles tab on the Diagnostics page of the Cloudera Management Console. This topic lists the diagnostic data that is generated.

Cloudera Data Services on premises diagnostics bundle

Namespace information (under <namespace>/):

- pods.json
- events.json
- replication.json
- services.json
- daemonset.json
- deployments.json
- replicaset.json
- configmap.json
- pv.json
- Pod logs (under logs/)
- Archive logs (under archived/)

bundle.json

- File name of diagnostics bundle
- Control plane version
- Control plane namespace

Environment information (under environment/)

- Per environment (under <environment-name>/)
 - Archive Logs (under archived/)
 - Non-Ozone logs (those stored in the persistent volume claim of the fluentd-aggregator pod)
 - Ozone logs (those stored in Ozone)
- cde
 - Separated into folders per namespace for the Cloudera Data Engineering services
 - Each folder contains namespace information as above
- dwx
 - Separated into folders per Cloudera Data Warehouse containing namespace information
 - DWX diagnostics bundle for each warehouse
- mlx
 - Separated into folders per MLX workspace
 - Each folder contains namespace information as above
- environment.json
 - Contains the following:
 - CM version,
 - CDH version
 - CM License UUID
 - Health of all of the hosts
- client-configs.zip
 - Client configs of the base cluster
- details.json

- Environment details

license.json

- License version
- License Name
- License UUID
- License Start date
- License Deactivation date
- License Expiration date

log.txt

- Good for debugging the diagnostics collection
- All log statements from the diagnostics collection are found here

monitoring-metrics.json

- The results of various monitoring queries

nodes.json

- Details of all cluster nodes
- Examples: labels, creationTimeStamp, etc.

Pod reaper logs (under pod-reaper/)


- Job logs for pod reaper

Collecting Cloudera on premises diagnostics data

You can collect and download Cloudera on premises diagnostic data for different components and services from the Diagnostics page of the Cloudera Management Console. In addition, you can specify the type of workloads and the duration for which you want to collect the diagnostic data.

Procedure

1. From the Cloudera Management Console home page, select Diagnostics. The Diagnostics homepage contains data related to the collected Bundles and Redaction Rules.



2. On the Diagnostics homepage, you can perform the following actions: **Collect New Bundle** and **Add New Rule**

Diagnostics

Bundles		Redaction Rules			
Search		Collect New Bundle			
☐	Status	Name	Size (MB)	Created At	
☐	Successful	cdp-diagbundle-b1120840-2026-08-18T04:59:43.275957535.zip	10.91	Aug 18, 2026 04:59 UTC	:
☐	Successful	cdp-diagbundle-b1120840-2026-08-18T04:48:14.137310297.zip	10.11	Aug 18, 2026 04:48 UTC	:
☐	Collecting	cdp-diagbundle-b1120840-2026-08-18T04:47:57.721350622.zip	0	Aug 18, 2026 04:47 UTC	:
☐	Successful	cdp-diagbundle-b1120840-2026-08-18T04:26:42.839644862.zip	9.34	Aug 18, 2026 04:26 UTC	:
☐	Collecting	cdp-diagbundle-b1120840-2026-08-18T04:06:54.175176622.zip	0	Aug 18, 2026 04:06 UTC	:
☐	Successful	cdp-diagbundle-b1120840-2026-08-18T03:57:20.200168047.zip	8.55	Aug 18, 2026 03:57 UTC	:

3. Select the Bundles tab and click Collect New Bundle. On the Collect New Bundle page, you can specify the following options for data collection: the time duration, the associated environments and services, and the size of the file to include the diagnostic data.

For more information on the diagnostic data options, see [Options for generating the Cloudera on premises diagnostic data](#) on page 4.

4. Click Collect and Send Diagnostic Data to commence the process.

Cloudera Management Console initiates the process of generating the diagnostic data based on the options that you select. After the data is generated, the Last Collected Time field is updated with the collection time.

Once the collection process is completed, navigate to the Diagnostics homepage and click on the collected bundle name to fetch more details. Alternatively, you can download or delete the selected bundle. The diagnostic data is downloaded in a .zip file. You can send this file to Cloudera support. If your system is connected to the internet, Cloudera Management Console collects the diagnostic data and sends it to Cloudera Support.

Bundles		Redaction Rules			
Search		Collect New Bundle			
☐	Status	Name	Size (MB)	Created At	
☐	Successful	cdp-diagbundle-a110df91-2026-08-20T08:24:02.641832094.zip	10.96	Aug 20, 2026 08:24 UTC	ⓘ

[Download](#)
[Delete](#)

Diagnostic Bundle Details

← Back Download Bundle Download Log Delete	
Details File Name cdp-diagbundle-b1120840-2026-08-18T04:26:42.839644862.zip Bundle Path /data-0/bundles/cdp-diagbundle-b1120840-2026-08-18T04:26:42.839644862.zip Log Path /data-0/logs/bundle-cdp-diagbundle-b1120840-2026-08-18T04:26:42.839644862.log Status Complete Size 9 MB Size Limit — Time Range (UTC) — Creation Time (UTC) 08/18/2026 4:26	Scope All Environments Progress - List Environments 1.0s - Collect Control Plane Kubernetes Data 51.2s - Collect Control Plane Logs 0.3s - Collect Analytical Metrics For DMX 0.3s - Collect Control Plane Monitoring Data 0.3s - Collect Global System Data 0.7s - Collect License Information 0.3s - Finalizing Bundle 1.2s - Upload Bundle 1.6s

5. On the Diagnostics home page select the Redaction Rules tab and click Add New Rule.

Diagnostics

Bundles Redaction Rules					
Define the list of rules to be followed for redacting sensitive information from the diagnostic bundle. Sensitive information, including but not limited to passwords and personal identifiers, is not collected.					
Search Add New Rule					
Description	Trigger	Search	Case-sensitive	Replace	Source
Redact passwords from...	password	'password'[!*"'\r\n]+	Disabled	'password': 'BUNDLE-REDACTED'	Built In
Redact password= an...	password	password[!*"'\r\n]+	Disabled	password=BUNDLE-REDACTED	Built In
Redact passwd= and ...	passwd	passwd[!*"'\r\n]+	Disabled	passwd=BUNDLE-REDACTED	Built In
Redact pass= and pass:	pass	pass[!*"'\r\n]+	Disabled	pass=BUNDLE-REDACTED	Built In
Redact PASSWORD,	PASSWORD,	PASSWORD[!*"'\r\n]+	Disabled	PASSWORD, BUNDLE-REDACTED	Built In
Redact key= and key:	key	key[!*"'\r\n]+	Disabled	key=BUNDLE-REDACTED	Built In
Redact secret= and secret:	secret	secret[!*"'\r\n]+	Disabled	secret=BUNDLE-REDACTED	Built In
Redact credential= an...	credential	credential[!*"'\r\n]+	Disabled	credential=BUNDLE-REDACTED	Built In
Redact token= and tok...	token	token[!*"'\r\n]+	Disabled	token=BUNDLE-REDACTED	Built In
Redact token= and tok...	token	token[!*"'\r\n]+	Disabled	token=BUNDLE-REDACTED	Built In

Redaction Rules define the list of rules to be followed for redacting sensitive information from the diagnostic bundle. Sensitive information, including but not limited to passwords and personal identifiers, is not collected.

New Redaction Rule ✕

When specifying a custom rule, make sure the Search field contains a regular expression to be matched against the data. If a match is found, it is replaced by the contents of the Replace field. Trigger is an optional simple string to enable faster searching by the regular expression. If the string is found, the redactor attempts to find a match for the regular expression. If no trigger is specified, redaction occurs by matching the regular expression. Use the Trigger field to enhance performance: simple string matching is faster than regular expression matching.

***Description**

A description for the redaction rule. Eg: Redact password= and password:

Trigger

An optional simple string to trigger faster searching by regular expression. Eg: password

***Search**

A regular expression that identifies the phrase to be replaced. Eg: password[!*"'\r\n]+

☐ **Case-sensitive**

A boolean value that indicates whether the search regular expression should be case-sensitive.

***Replace**

A string that replaces the phrase to be redacted. Eg: password=BUNDLE-REDACTED

Add New Rule Cancel

6. Input the information for the following:

- **Description:** (Mandatory) Enter the description for the redaction rule.
- **Trigger:** Enter the redaction rule trigger string.
- **Search:** (Mandatory) Enter the redaction rule search expression.

- Replace: (Mandatory) Enter the redaction rule replace string.
7. Click Add New Rule. The defined redaction rule is created.
 8. Click Cancel to return to the previous page.